arXiv:1608.08953v2 [cs.HC] 9 Feb 2017

Dynamic Allocation of Crowd Contributions for Sentiment Analysis during the 2016 U.S. Presidential Election

Mehrnoosh Sameki, Mattia Gentil, Kate K. Mays, Lei Guo, and Margrit Betke*

Boston University

Abstract

Opinions about the 2016 U.S. Presidential Candidates have been expressed in millions of tweets that are challenging to analyze automatically. Crowdsourcing the analysis of political tweets effectively is also difficult, due to large inter-rater disagreements when sarcasm is involved. Each tweet is typically analyzed by a *fixed* number of workers and majority voting. We here propose a crowdsourcing framework that instead uses a *dynamic* allocation of the number of workers. We explore two dynamic-allocation methods: (1) The number of workers queried to label a tweet is computed *offline* based on the predicted difficulty of discerning the sentiment of a particular tweet. (2) The number of crowd workers is determined *online,* during an iterative crowd sourcing process, based on inter-rater agreements between labels. We applied our approach to 1,000 twitter messages about the four U.S. presidential candidates Clinton, Cruz, Sanders, and Trump, collected during February 2016. We implemented the two proposed methods using decision trees that allocate more crowd efforts to tweets predicted to be sarcastic. We show that our framework outperforms the traditional static allocation scheme. It collects opinion labels from the crowd at a much lower cost while maintaining labeling accuracy.

1 Introduction

During the 2016 U.S. presidential primary election season, the political debate on Twitter about the four presidential candidates Hillary Clinton, Ted Cruz, Bernie Sanders, and Donald Trump was particularly lively and created a huge corpus of data. It has been argued that Twitter can be considered a valid indicator of political opinion (Tumasjan et al. 2010), and so various parties, including journalists, campaign managers, politicians, and social scientists, are interested in using automated natural language processing tools to mine this corpus.

Unsupervised learning methods have been used previously to analyze a similar corpus, 77 millions tweets about the 2012 U.S. presidential election and create summary statistics such as "twitter users mentioned foreign affairs in connection with Obama more than with Romney" (Guo et al. 2016). Supervised learning methods also have been used, for example, to analyze filtered "snippets" of political blogs (Hsueh, Melville, and Sindhwani 2009). However, creating *accurate* learning methods to analyze positive or negative sentiments is challenging. Political opinions expressed on the internet often contain sarcasm and mockery (Guo et al. 2016; Hsueh, Melville, and Sindhwani 2009), which are difficult to discern by machine or human computation (González-Ibáñez, Muresan, and Wacholder 2011; Young and Soroka 2012)

Crowdsourcing has been proposed to collect training data for predictive models used to classify political sentiments (Hsueh, Melville, and Sindhwani 2009; Wang et al. 2012). Out of concern for the accuracy of human annotation, it is standard practice to collect multiple labels for the same data point and then use the label that obtained a majority vote (Karger, Oh, and Shah 2013). Typically an odd number of crowd workers, e.g., five or seven, is chosen to create this redundancy. Redundancy, however, cannot guarantee reliability, i.e., agreement among the raters with each other about the sentiment present in the text in question. For example, when five crowd workers analyzed the sentiments expressed in the political snippets dataset (Hsueh, Melville, and Sindhwani 2009), only a 47% agreement rate on the three labels "positive," "negative," or "neutral sentiment" could be achieved.

Hsueh et al., 2009, noted that "not all snippets [of political blogs] are equally easy to annotate." We made the same observation for our data – sarcastic twitter messages are more difficult to label, and we propose to allocate crowd resources according to the predicted difficulty level: The more difficult the sentiment analysis may be, the higher the number of workers becomes that our model assigns. In allocating fewer crowd workers to tasks that are predicted to be easy, we aim to balance the goals of labeling accuracy and efficiency.

The literature describes techniques for optimal tradeoffs between accuracy and redundancy in crowdsourcing (Karger, Oh, and Shah 2013; Tran-Thanh et al. 2013). In these works, the proposed crowdsourcing

* Presented at HCOMP 2016
Corresponding Author: M. Betke, betke@bu.edu.
Copyright © 2016. Sameki, Gentil, Mays, Guo and Betke. All rights reserved.

mechanism uses a fixed number of crowd workers per task, and the assignment is agnostic about the latent difficulty level of each task. If the difficulty of a task can be discerned, easy tasks could be routed to novice workers and difficult tasks to expert annotators (Kolobov, Mausam, and Weld 2013). Optimal task routing, however, is an NP-hard problem, and so online schemes for task-to-worker assignments have been proposed (Bragg et al. 2014; Rajpal, Goel, and Mausam 2015). Our work falls into this category of online crowdsourcing methodology.

Our contributions are as follows:

- We propose a decision-tree approach for *dynamically* determining the number of crowd workers for tasks that require redundant annotations.
- We provide two versions of this approach: The *offline* version computes the number of workers needed based on the content of the data they are asked to analyze. The *online* version relies on iterative rounds of crowdsourcing and determines the number based on content and annotation results in previous rounds.
- To illustrate and evaluate our approach, we conducted a crowdsourcing experiment with a dataset of 1,000 tweets that were sent during the 2016 primary election season. We collected 5,075 ratings of the sentiment towards presidential candidates Clinton, Cruz, Sanders, and Trump in these tweets and evaluated their accuracy with respect to a gold standard established by experts in political communication.
- Comparisons with traditional crowdsourcing strategies show that the proposed offline and online selection methods intelligently detect ambiguities in sentiment analysis and recruit more workers to resolve those. We show that a large portion of the crowdsourcing budget can be saved at a small loss of accuracy.

2 Method

We here describe our method to solve the problem of dynamically assigning crowd workers to analyze the sentiment of political tweets. Our approach consists of three main components. First, we designed a method to detect sarcasm in tweets (Section 2.1). This first step was important because sarcasm is one of the most confusing and misleading language features to classify even for a human annotator, especially when a single out-of-context tweet is being analyzed.

We then constructed a decision tree that assigns to each tweet a fixed number of crowd workers based on the presidential candidates mentioned in the tweet and other text properties, in particular, its sarcasm (Section 2.2). In designing such a tree, we were motivated by the following insight: For tweets which are expected to be clear and straight-forward to analyze, fewer annotators would be required than for tweets that are sarcastic and complicated. To build the tree, we estimated how troublesome it would be for a crowd worker to correctly understand what kind of sentiment is being expressed towards the candidates.

The third component of our approach moves from an offline to an online consideration of how many crowd workers to involve in the labeling process (Section 2.3). Based on the inter-rater agreements between labels obtained in a first phase of an iterative crowd sourcing process, for tweets which proved to be challenging to annotate, our method determines how many additional labels to acquire in one or more subsequent crowd sourcing phases.

Our final methodological contribution is a description of the equivalency between two crowdsourcing schemes, the traditional 5-worker-per-task scheme and the dynamic scheme that assigns 3 workers per task in the first round and 2 additional workers in a second round if disagreement is encountered in the 1st round. This is a general result about offline versus online crowdsourcing schemes. It holds for any application and is therefore presented in Section 2.4, separate from the results of our sentiment analysis of political tweets.

2.1 Sarcasm Detection

Our first step was trying to predict whether a given tweet was sarcastic or not. We used a Bayesian approach to estimate the likelihood of sarcasm based on training data provided by domain experts. Our training data contains the label "sarcasm present" or "sarcasm not present" for 800 tweets about the four presidential candidates Clinton, Cruz, Sanders, and Trump.

We looked for general features that are usually clues for the presence of sarcasm in a sentence (González-Ibáñez, Muresan, and Wacholder 2011; Davidov, Tsur, and Rappoport 2010) and grouped them into 7 categories:

1. Quotes: People often copy a candidate's words to make fun of them.
2. Question marks, exclamation or suspension points.
3. All capital letters: Tweeters sometimes highlight sarcasm by writing words or whole sentences with all-capital letters.
4. Emoticons like ':)', ':('
5. Words expressing a laugh, or other texting lingo, such as 'ahah,' 'lol,' 'rofl,' 'OMG,' 'eww,' etc.
6. The words 'yet' and 'sudden.'
7. Comparisons: Many tweeters use comparisons to make fun of a candidate, using words such as 'like' and 'would'.

The sarcasm detecting algorithm that we designed scans the tweet text for those features and returns the list of sarcastic clues. The clues are represented by a 7-component feature vector f that contains a Boolean value for each of the categories listed above – "1" indicates "presence" of the feature, "0" otherwise.

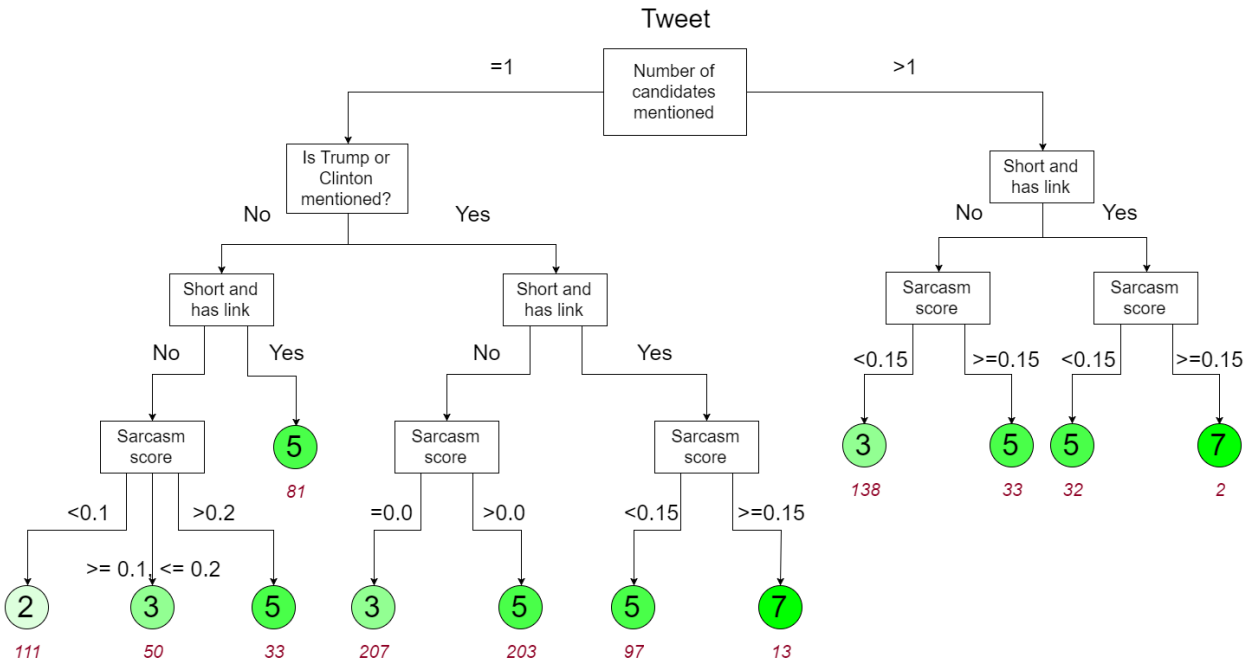

Figure 1: The Static Decision Tree (SDT) model used to determine the number of crowd workers (leaves) to engage in analyzing tweets about four presidential election candidates. The intensity of the leaf shading visualizes costs, e.g. pale green corresponds to low costs. The sarcasm score is computed according to Eq. 5. Experimental results are shown under each leaf as the number of tweets processed (red).

Given a tweet t and its feature vector f, our method computes the probability that the tweet t contains sarcasm by using Bayes rule:

$$P(t \text{ is sarcastic} | f_n) = \quad (1)$$

$$\frac{P(f_n | t \text{ is sarcastic})\, P(t \text{ is sarcastic})}{P(f_n)} = \quad (2)$$

$$\frac{\#\text{ of sarcastic tweets with } f_n}{\#\text{ of tweets with feature } f_n}. \quad (3)$$

To weigh the presence of the n-th feature in sarcastic tweets appropriately, our method computes a weight vector w by normalizing its n-th component by the probability that it is sarcastic, given any of the seven features is present:

$$w_n = \frac{P(t \text{ is sarcastic} | f_n)}{\sum_{n=1}^{7} P(t \text{ is sarcastic} | f_n)}. \quad (4)$$

Our sarcasm score for each tweet is then defined to be the dot product

$$w^T f \quad (5)$$

of the weight and feature vectors.

2.2 Decision Tree

The decision tree we designed maps a tweet to a number of crowd workers that will be asked to label the tweet. To gain insight into the properties of a tweet that could cause a crowd worker to struggle in sentiment classification and warrant additional crowd work, we obtained gold standard data and conducted a formative crowdsourcing study.

Expert Labels We used 1,000 tweets about the four presidential candidates Clinton, Cruz, Sanders, and Trump. For these tweets, we had gold standard labels about two categories, provided by experts in political communication. The first category was whether each of the four candidates was mentioned in the tweet. The second category described whether the tweet was in general "positive," "neutral" or "negative" about each candidate mentioned in the tweet. If more than one candidate was mentioned in a tweet, the sentiment towards each candidate was labeled.

Formative Crowdsouring Experiment We asked 5 crowd workers to analyze each tweet, calling our experiment the *"Trad 5 baseline"* (the details on the crowdsourcing methodology are given in Section 3). We asked the workers who among the four candidates Sanders, Trump, Clinton and Cruz was mentioned and to indicate the attitude that the tweeter expressed towards them on a three-point scale "positive," "neutral," or "negative."

Decision Tree Design We designed our decision tree (see Fig. 1) based on the properties we observed that influence the accuracy with which a worker interprets the sentiment of the tweet. The first branching of the tree

accounts for whether one or more candidates are mentioned in the tweet text, the most relevant factor in its sentiment analysis. Tweets in which several candidates are mentioned are more difficult to classify because annotators can become confused by the different attitudes that the writer expresses towards each of the candidates or by the presence of comparisons between them. We here provide three examples:
Tweet 1

> *@BecketAdams @JPTruss @GayPatriot except Cruz now realises Trump's power and is debating him. Rubio is still hiding from Trump on stage*

is "positive" towards Trump and "neutral" towards Cruz, according to expert opinion. Four crowd workers agreed that the message was "neutral" towards both candidates, and one labeled it "positive" towards Trump and "neutral" towards Cruz.
Tweet 2

> *Bernie's Super PAC Hypocrisy: Twice as Much Outside Money Spent Supporting Sanders as Promoting Clinton https://t.co/RVAi7X4shS*

is "positive" towards Clinton and "negative" towards Sanders, according to expert opinion. All five crowd workers agreed but not on the correct labels – they selected a negative sentiment towards Sanders and a neutral for Clinton.
Tweet 3

> *Has Trump mentioned that he doesn't think Cruz is eligible to be President recently? That seemed like a go-to for him*

misled annotators because both sarcasm is present and two candidates are mentioned. As a consequence, only 3 workers out of 5 agreed on a negative overall feeling towards both candidates.

It is important whether Clinton or Trump was mentioned in the Tweet. Opinions towards these candidates are usually more challenging to understand as tweeters have very disparate and unclear attitudes towards them.

The next layer of the decision tree accounts for the length of the tweet and the presence of a link. We consider a tweet short if it contains fewer than 10 proper words. Tweets that contain a webpage address are not always fully understandable by themselves as they refer to the content of the link or they are a response to another tweet, and therefore their context is not always clear.

Finally, the terminating decision layer in the tree is based on the sarcastic score that was produced by the sarcasm predictor. The decision tree uses the sarcasm score as defined in Eq. 5 to determine the likelihood of sarcasm in the particular tweet.

We assigned a fixed number of crowd workers to each leaf of the tree, which specifies the number of annotations needed for a particular tweet. In this first model we grouped the tweets into 4 categories (very easy, easy, medium and hard) and assigned 2, 3, 5, or 7 workers to them respectively. We call the model "Static Decision Tree" (SDT) due the fact that the number of crowd workers depends only on the content analysis of the tweet (and not dynamically on the workers' labels, as described below). With this tree, the number of crowd workers to be queried for each tweet can be computed *offline* – in advance of any crowdsourcing experiment (i.e., the numbers shown in Fig. 1 with a green-shaded background).

2.3 Dynamic Worker Assignment

We here propose an *online* scheme for determining the number of crowd workers to be queried for each tweet. This approach cannot be computed in advance to the crowdsourcing experiment but is an iterative method that relies on the results of the crowd work.

Our idea is to request a low number of workers to provide the sentiment analysis of each tweet in a first round of crowdsourcing, and then perform one or more rounds of crowdsourcing for the tweets for which workers disagreed. In this way, the difficulty of the tweet is observed directly as a measure of disagreement in the first round of crowdsourcing, and we do not risk wasting effort on tweets that are trivial to classify. To evaluate our approach, we designed two instantiations of our idea involving two rounds of crowdsourcing:

Dynamic Decision Tree 1 (DDT1) The first dynamic tree assigns 2 workers to the 'very easy' and 'easy' difficulty classes, 3 for 'medium' and 5 for 'hard.' If the 2 workers disagree on classifying a 'very easy' or 'easy' tweet, we conduct a second round of crowdsourcing on that tweet so that we can get a majority vote. If some annotators disagree for a 'medium'-class tweet, 2 more workers are involved. The number of workers for 'hard' tweets stays fixed.

Dynamic Decision Tree 1 (DDT2) Finally, we pushed the dynamic assignment design even further and set up a tree that starts with a very low numbers of annotators in order to minimize the number of crowdsourced tasks. This tree initially assigns 2 workers to the 'very easy' and 'easy' classes and requires 3 more annotators if the initial workers disagree. The tweets in the 'medium' and 'hard' categories were first only analyzed by 3 workers, and this number is increased by 2 workers if at least one disagreement is observed.

2.4 Equivalency of Traditional Static versus Proposed Dynamic Worker Allocation

Past work showed that the probability p that a crowd worker w correctly performs a task t according to a gold standard label can be described as a function $p(t, w)$ of the task difficulty and the worker skill (Ho and Vaughan 2012). For simplicity of our analysis, we omit the dependence on the worker.

For a generic task, we can compute the probability P_M that the gold standard is successfully obtained by majority voting for a set of crowd sourcing baseline

schemes as a function of p. For example, the probability P_M that the traditional 3-worker-per-task crowdsourcing scheme yields the correct results is the probability that at least 2 out of 3 performed the task correctly, which is

$$P_M = \sum_{i=2}^{3} P(i \text{ workers are correct}) = \sum_{i=2}^{3} \binom{3}{i} p^i (1-p)^{(3-i)} = p^2[3(1-p)+p]. \quad (6)$$

Similarly, with the traditional 5-worker-per-task crowdsourcing scheme, we attain $P_M =$

$$\sum_{i=3}^{5} P(i \text{ workers are correct}) = \sum_{i=3}^{5} \binom{5}{i} p^i (1-p)^{(5-i)} = p^3[10(1-p)^2 + 5p(1-p) + p^2]. \quad (7)$$

Next we simulate the dynamic assignment of workers with 3 initial workers, where 2 more workers are involved if disagreement is encountered. The probability that this model produces the correct result by majority voting is the sum of three probabilities: (1) the probability that the three initial workers agree on the correct result, (2) the probability that one initial worker performs the task incorrectly and at least one new worker correctly, and (3) the probability that only one initial worker performs the task correctly and both the new workers follow up correctly:

$$\binom{3}{3} p^3 + \left[\binom{3}{2} p^2(1-p)\right](1-(1-p)^2) + \left[\binom{3}{1} p(1-p)^2\right] p^2 = p^3[1+3(1-p)(2-p)+3(1-p)^2)] = p^3[10(1-p)^2 + 5p(1-p) + p^2]. \quad (8)$$

The derivations in Eqs. 7 and 8 result in the same formula. We can therefore infer that a dynamic 3(+2) allocation method for workers achieves the same prediction accuracy as the traditional 5-worker crowdsourcing scheme. As we will describe in more details below, by running such a model on all tweets in our dataset we were able to obtain optimal results from crowdsourcing with only 4,058 tasks. This result is impressive because it proves that we can reach exactly the same accuracy level and save 18.84% of our budget only by running two "smart" rounds of crowdsourcing.

3 Experimental Methodology

Our data consists of 1,000 tweets about the four presidential candidates Clinton, Cruz, Sanders, and Trump sent during the primary election season in February 2016. We selected these candidates because they were the two leading candidates in the polls at the time of data collection from each major U.S. political party (Republican and Democrat). The data were collected by using the Crimson Hexagon ForSight social media analytics platform (http://www.crimsonhexagon.com/platform).

The tweets were labeled by two domain experts with a background in political communication in a two-phase process. In the first phase, the experts determined the sentiment towards each candidate mentioned in each tweet independently. In the second phase, they came to a consensus on the tweets that they had initially disagreed on.

For our crowdsourcing experiments, we used the Amazon Mechanical Turk (AMT) Internet marketplace to recruit workers. We accepted all workers from the U.S. who had previously completed 100 HITs and maintained at least a 92% approval rating. We paid each worker \$0.05 per completed task. We conducted two crowdsourcing studies, a formative and a summative study, involving 200 and 800 tweets respectively.

Formative Study. We gave the following instruction before presenting every tweet:

> *Carefully read through each tweet and decide the author's attitude toward each mentioned presidential candidate (support, neutral, or against).*

We verified that short tweets (fewer than 10 proper words) were very difficult to tag. Tweets with links to an external page were also difficult to analyze. It is likely that the sentiment of the tweet heavily relies on the content of the referenced webpage. Workers may have tried to follow the link or may have selected a random sentiment instead of following the link. In our instructions for our summative study, we therefore specifically asked the crowd workers not to click on any external link for completing the task. We also adjusted the label for positive and negative sentiments towards a candidate.

Summative Study. We updated the instructions as follows:

> *Read through the tweet and answer the following questions. Do NOT click on any links.*
>
> *Read the tweet and decide whether the candidate was mentioned at all or not. Note that the reference of Twitter user names (e.g., @realDonaldTrump, @HillaryClinton) or hashtags (e.g.,#Trump2016, #HillaryClinton2016) is also counted as a mention.*
>
> *Express which sentiment was manifested by the writer towards them: positive, neutral, or negative.*

We collected ratings from a traditional crowdsourcing scheme that involves 5 independent workers per tweet. We call this the "Trad 5" baseline. For 15 tweets that were deemed 'hard' to analyze by our decision tree and thus required the ratings from 7 workers, we needed to collect additional ratings. Instead of simply collecting two more, we asked for 5 additional ratings per tweet from which we could then draw additional samples randomly for analysis. This resulted in a total of 5,075 tasks.

To simulate a crowd sourcing experiment that employs a fixed number of three crowd workers per tweet (our traditional Trad 3 baseline), we randomly sample the results produced by 5 crowd workers. To simulate the crowd sourcing experiments that use the decision tree we designed (SDT, DDT1, DDT2), we similarly use random samples from our Trad 5 baseline. To obtain the results of our decision trees, we averaged the collected metrics over 5 different model runs to attenuate potential noise generated by the randomness in selecting crowd workers.

Evaluation Measures We use two metrics for evaluating our work. They are meaningful for understanding the trade-off between accuracy and budget concerns, which is the focus of our work.

- **Number of crowd worker tasks:** This is the total number of Human Intelligence Tasks requested by our decision tree model. The number provides an indication of the budget needs of a crowd experiment. To find the monetary costs of crowdsourcing, we can multiply this number by the price per task (we used \$0.05/task).
- **Accuracy of the labeling:** The accuracy of the crowdsourced sentiment analysis can be determined by how much agreement exists between the majority crowdsourced opinion and the gold standard opinion provided by experts. Our main measure of accuracy is Cohen's Kappa score κ for measuring inter-rater reliability (IRR). Cohen's Kappa score accounts for the possibility that raters are guessing and so an agreement is obtained by chance.

4 Results

Sarcasm detection Our experiments showed that the clues we used for sarcasm detection are very diverse, and were used in different ways according to the topic of the tweet. We found that smileys were not used at all, while the most meaningful element for sarcasm detection was the presence of expressions like 'lol', 'hahaha,' for example, in the following tweet:

> *RT @rickygervais If Trump was a teacher he'd be fired for publicly saying the things he says. Luckily he isn't a teacher. Just the next president. Hahaha.*

The presence of sarcasm was indeed a factor which increased the difficulty of tweet classification: in our dataset, sarcastic tweets had a 71.2% percentage inter-rater agreement. This metric increased to 78.3% when dealing with non-sarcastic tweets.

It turned out that the presence of sarcasm was not as ubiquitous as we had expected, as only 73 messages out of 800 were estimated to be sarcastic by domain expert, and a surprising 68.5% of them concern Donald Trump (see Table 1). The last row of the table shows that even after weighing the sarcasm presence over the number of tweets that mentioned each candidate, Donald Trump still leads with 12% of his tweets that are sarcastic. Regarding the sentiment that is usually associated with sarcasm, the last column of the table proves that sarcasm is usually associated with a negative feeling towards a candidate. In fact this language feature is usually employed to make fun of a candidate and criticize him for his statements or actions.

	Clinton	**Cruz**	**Sanders**	**Trump**	
Positive	0	0	2	3	5
Neutral	5	6	4	15	30
Negative	6	4	6	32	48
Sarcastic	11 5.9%	10 6.2%	12 6.8%	50 12.0 %	

Table 1: These results show the number of sarcastic tweets addressed to each candidate and the sentiment that they showed according to the gold standard provided by experts in political communication. In the dataset of 800 tweets, 73 tweets were sarcastic. The last row shows the ratio of sarcastic tweets over the total tweets in which each candidate was mentioned.

Differences Based on Specific Candidates As expected, we found that which presidential candidate was mentioned in a tweet had an impact on how difficult it was to discern the tweeter's opinion about the candidate. The sentiments that tweeters expressed towards Hillary Clinton and Donald Trump were often unclear or veiled by sarcasm. To illustrate this point qualitatively, we give an example tweet about Trump that confused the crowd workers:

> *I was watching the Texas gop debate on snapchat lol and this is the only state where I've seen people actually rally against trump YOUNG PPL.*

One crowd worker labeled the tweet to show "a positive attitude," 2 crowd workers labeled it as "neutral" and the remaining 2 agreed on a "negative" sentiment towards the candidate. In this case, it is impossible to determine a result by majority vote, and a final label can be assigned by a reasonable random choice. We here chose randomly between "neutral" and "negative."

To illustrate the issue quantitatively, we here provide the inter-rater reliability values among 5 crowd workers of our formative study when classifying sentiments towards each candidate and report both the relative observed agreement among crowd workers and Cohen's Kappa score κ:

Candidate	Agreement	Kappa IRR
Bernie Sanders:	83.05%	$\kappa = 0.74$
Ted Cruz:	87.78%	$\kappa = 0.78$
Hillary Clinton:	63.41%	$\kappa = 0.41$
Donald Trump	78.13%	$\kappa = 0.66$

It is evident from the above numbers that annotators disagreed much more often when Clinton or Trump were mentioned. For our summative study, we therefore designed an offline model that can account for this ob-

	Trad 3	Trad 5	SDT	DDT1	DDT2
Efficiency	3,000	5,000	3,907	3,206	3,608
Imprv.			22%	36%	28%
Accuracy	0.612	0.653	0.624	0.630	0.643
Loss			4.4 pp	3.5 pp	1.0 pp

Table 2: Comparison of results of five methods with respect to their efficiency and accuracy. The number of crowd workers engaged (i.e., efficiency or costs) and the accuracy of their sentiment labeling (Cohen's Kappa IRR rate) compared to the gold standard established by experts are given for each method. For the first two methods, each tweet is analyzed by the same fixed number of crowd workers, i.e., 3 crowd workers (Trad 3) or 5 crowd workers (Trad 5). For the methods that use the decision tree (DT), the number of crowd workers engaged depends on the content of the tweet and result in significant improvements (Improv.) in efficiency with respect to the 5 crowd-worker models (row 2), without much loss of accuracy (row 4, given in percent points, pp).

servation and involve more workers to label tweets from these two candidates.

Results for Traditional Fixed-Allocation Model The first two models that we considered are a fixed crowdsourcing round with the same amount of workers for every tweet. With a total of 3 annotators we requested 3,000 ratings and we achieved a 0.612 Kappa value (see Table 2). If we increase the number of crowd workers by 2 we require 5,000 tasks and we would get a 0.653 reliability measure. These results align with previous observations that the task of sentiment analysis is challenging even for human annotators (Young and Soroka 2012; Tumasjan et al. 2010) Despite the significantly higher costs of requesting 2,000 additional labels from crowd workers, a 40% increase, the average agreement between the majority of crowd contributions and expert labels improved by only 6.3 percent (or, equivalently, by a difference of Kappa values of 4.1 percent points).

Results for the Proposed Static Decision Tree For the static decision tree (SDT), 3,907 labels were requested, on average, and an IRR score of 0.624 was obtained. The allocated numbers of workers based on the text analysis of the tweets and decision rules of the tree are shown in red in Figure 1. With this static decision tree, 22% of the budget would be saved with respect to the traditional 5-worker-per-task model (Trad 5). The loss in accuracy is 4.4 percent points.

Results for the Proposed Dynamic Decision Trees The first dynamic tree (DDT1) showed a meaningful improvement as it involves only 3,206 tasks on average and has an IRR score of 0.630. This model costs 36% less than the fixed one with 5 workers and only 6.9% more than the model with 3 annotators but the gain in accuracy with respect to the latter is quite high (2.9%). This model would be preferable in low-budget situations.

The second dynamic tree (DDT2) is a bit more expensive as it requires 3,608 annotators by average but the Cohen's Kappa IRR rate improves to 0.643. Even this classifier is much cheaper than the fixed 5-worker as it saves almost 28% of the budget and the accuracy is comparable (the difference between Kappas scores is only 1 percent point). We propose that this predictor is suitable if we are willing to spend a bit more in order to achieve a very good performance.

Both dynamic trees produce notably better results than the fixed decision tree in both cost and accuracy. This shows that the difficulty of a tweet can be inferred from the crowdsourcing outcomes themselves and that heuristic rules for determining it are extremely complex and hard to formulate. Correct results can be obtained by a second round of annotations, which needs to be set up accordingly, thus saving a meaningful amount of budget.

Cost Savings of Dynamic versus Static Worker Assignment The traditional 5-worker-per-task allocation model Trad 5 performs exactly the same as a dynamic model which assigns 3 annotators +2 more if there is disagreement, as described in Section 2.4. This result shows that our model allows the same accuracy but at a much lower cost. A visualization of the differences in accuracy and efficiency between traditional static crowdsourcing schemes and the proposed dynamic schemes is given in Figure 2.

Analysis of Crowd Work Properties We submitted 5,075 tasks to Mechanical Turk for an overall cost of $253.75. The number of MT workers who contributed labels to all the tweets was 218. An average of 23 annotations was submitted per worker.

We analyzed how much time workers spent in labeling a single tweet, which is illustrated in Figure 3. Annotators spent an average of 85.1 seconds for classifying a single message but some workers were very meticulous and used up to 10 minutes to complete a single task. For example one of the best annotators who worked for us labeled 217 tweets with an average of 212 seconds per task, which sums up to almost 13 hours spent on the platform. On the other hand, other annotators were very quick, for instance one worker contributed by labeling 42 tweets and spent on average less than 9 seconds per message.

Sample Results on Political Tweets Analysis of the annotations of our 1,000 tweet dataset provides some fascinating observations about political opinions. We can report the overall sentiment that people showed towards candidates, as rated by the crowd workers (Table 3) and by the experts in political communication (Table 4). We found that Trump is the "most popular" candidate to tweet about, considering that more than half of the total tweets mentioned him, while the other candidates were evenly referred to on average. Further-

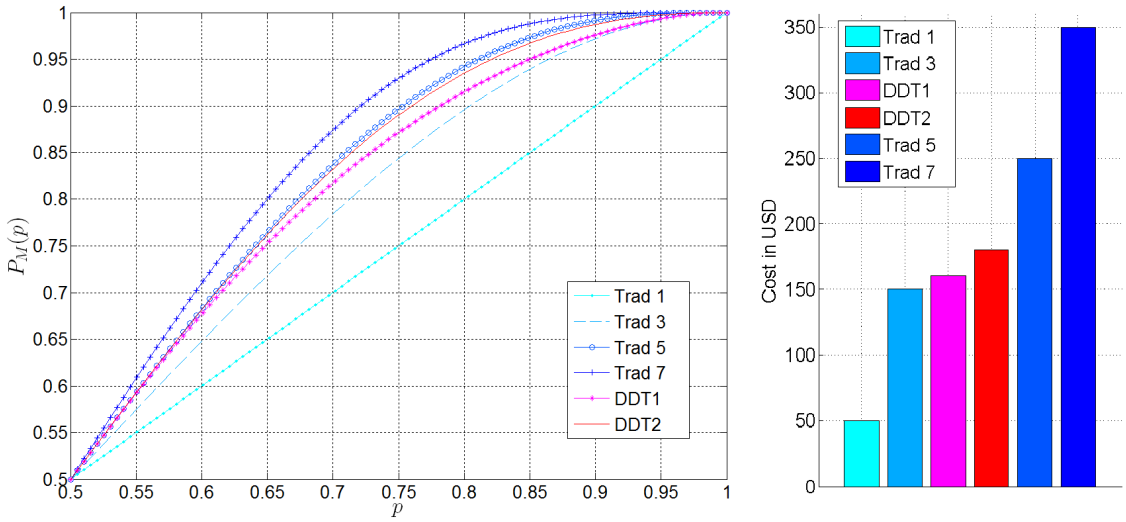

Figure 2: Performance Analysis – Accuracy and Costs. Left: The probability $P_M(p)$ that a given crowdsourcing scheme produces the correct label by majority vote as a function of the probability that a certain tweet is labeled correctly by a worker. We compare the performance of four traditional crowdsourcing baselines (with 1, 3, 5 or 7 crowd workers for each tweet) and our dynamic prediction models DDT1 and DDT2. For tweets that are easy to annotate, the accuracy of all methods is similar. When tweets are more difficult to analyze, and thus, more workers are engaged, the performance gains in accuracy of the DDT1 and DDT2 models compared to the traditional models "Trad 3" become apparent. The DDT2 model almost reaches the performance of the baseline "Trad 5." Right: The proposed dynamic models DDT1 and DDT2 provide large budget savings.

	Clinton	**Cruz**	**Sanders**	**Trump**	
Positive	33	44	78	123	278
Neutral	86	67	104	189	446
Negative	99	99	56	201	455
	218	210	238	513	

Table 3: Number of tweets, out of a total of 800, grouped according to crowd-sourced sentiment label per candidate. The last row and columns display the sums over the columns and rows respectively in the table.

	Clinton	**Cruz**	**Sanders**	**Trump**	
Positive	25	37	58	90	210
Neutral	109	85	123	208	525
Negative	91	89	55	212	447
	225	211	236	510	

Table 4: Number of tweets, out of a total of 800, grouped according to expert-provided sentiment label per candidate. The last row and columns display the sums over the columns and rows respectively in the table.

more it is clear that tweeters who discuss candidates for presidential elections often express negative feelings and complain about candidates, since there are about twice as many negative messages than positive ones in our entire dataset. The main difference between the crowd worker and expert annotations was the tendency of the crowd worker to label fewer tweets as "neutral."

5 Discussion and Conclusions

As crowdsourcing becomes more and more popular for large scale information retrieval, the cost of this human computation is becoming relevant. Example applications are real-time sentiment analysis to provide fast indications of changes in public opinion or collection of a sufficiently large training data for machine learning methods for big data analytics (Wang et al. 2012). Investigations, as ours, about how to balance the goals of efficiency and accuracy in crowdsourcing, are therefore particularly timely.

Few works have explored dynamic approaches to crowdsourcing that rely on iterative rounds of crowdsourcing and determine the number of worker assignments based on content and annotation results in previous rounds (Bragg et al. 2014; Ho and Vaughan 2012; Kolobov, Mausam, and Weld 2013). Connections to active and reactive learning (Yan et al. 2011; Lin, Mausam, and Weld 2015) have been made. While prior work involves theoretical analysis and simulation studies, we here provide a concrete solution to the problem of analyzing the sentiment of political twitter messages using a dynamic worker allocation framework.

We proposed a dynamic two-round crowdsourcing scheme that we embedded into a decision tree classifier. Other types of classifiers may be used, and, in future work, we will explore additional learning methods.

Analysis of political tweets is challenging due to the

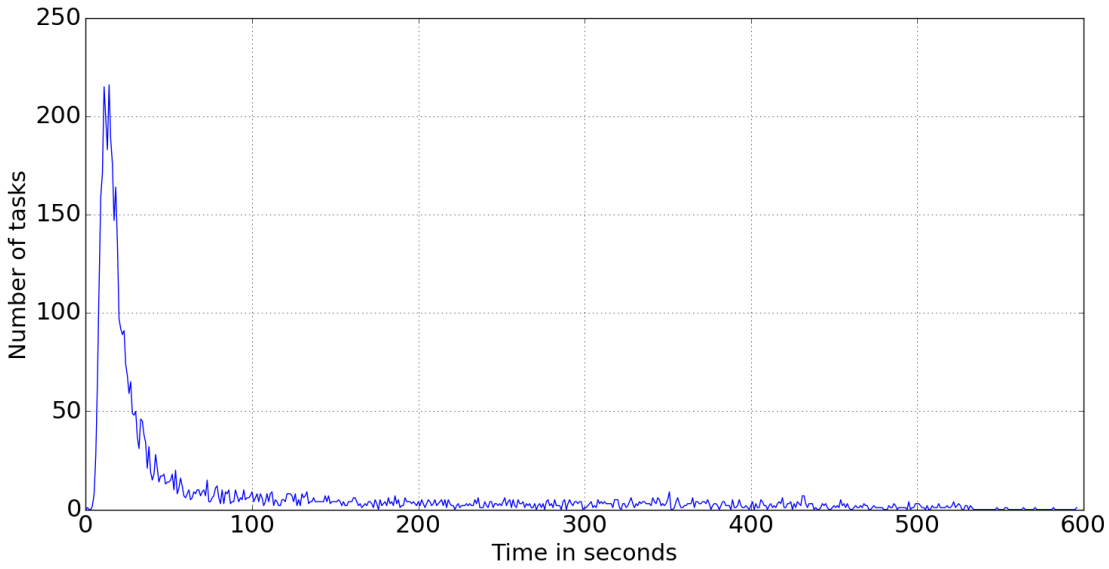

Figure 3: A distribution of tasks (HITs) as a function of task time, ranging from 1 to 600 seconds. This distribution was computed over the total 5,075 tasks that were submitted to Amazon Mechanical Turk during our crowdsourcing experiment.

short text and unknown context. Sentiment analysis is particularly difficult. Existing off-the-shelf text analysis systems can only provide a single sentiment label for a given text automatically. We found that they fail to distinguish the separate sentiments that were expressed when more than one presidential candidate was mentioned in a tweet. The presence of sarcasm exacerbated the problem. Our proposed solution is to design a classifier that early in the analysis makes a decision about the number of sentiments that must be revealed. Our new dataset may inspire other researchers to develop text analysis tools that address the difficult problem of multi-sentiment analysis and sarcasm detection.

Our corpus of 1,000 twitter messages is unique because it includes information about (1) the presence/absence of sarcasm and (2) a label about the specific sentiment for each candidate mentioned in the tweet (positive, neutral, negative), as determined by consensus of two domain experts.

It is notable that our study involved communication researchers in many aspects of the research, such as the development and refinement of crowdsourcing task instructions and the design of the Mechanical Turk interface. The intervention of domain experts greatly helped improve the validity and performance of our crowdsourcing method.

Likewise, the proposed approach has the potential to make a significant contribution to communication research. Traditionally, communication researchers use manual content analysis, a method that usually relies on two or three human coders, to analyze text in different media outlets or that of public opinion (Riffe, Lacy, and Fico 2014). However, the traditional method is tedious, time consuming, and limited by the nature of human subjectivity. Arguably, the use of the dynamic online crowdsourcing framework introduced in this study allows communication researchers to process larger datasets in a more efficient and reliable manner. Given the results of the study, future research should also consider cross-disciplinary collaboration to advance theories and methods for large-scale text analysis.

Acknowledgments

The authors would like to thank the Boston University Rafik B. Hariri Institute for Computing and Computational Science and Engineering for financial support and the crowd workers for their annotations.

References

[Bragg et al. 2014] Bragg, J.; Kolobov, A.; Mausam, M.; and Weld, D. S. 2014. Parallel task routing for crowdsourcing. In *Second AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2014)*, 11–21.

[Davidov, Tsur, and Rappoport 2010] Davidov, D.; Tsur, O.; and Rappoport, A. 2010. Semi-supervised recognition of sarcastic sentences in twitter and amazon. In *Proceedings of the Fourteenth Conference on Computational Natural Language Learning*, CoNLL '10, 107–116.

[González-Ibáñez, Muresan, and Wacholder 2011] González-Ibáñez, R.; Muresan, S.; and Wacholder, N. 2011. Identifying sarcasm in twitter: A closer look. In *Proceedings of the 49th Annual Meeting of*

the Association for Computational Linguistics: Human Language Technologies, 581–586.

[Guo et al. 2016] Guo, L.; Vargo, C. J.; Pan, Z.; Ding, W.; and Ishwar, P. 2016. Big social data analytics in journalism and mass communication: Comparing dictionary-based text analysis and unsupervised topic modeling. *Journalism and Mass Communication Quarterly* 1–28.

[Ho and Vaughan 2012] Ho, C.-J., and Vaughan, J. W. 2012. Online task assignment in crowdsourcing markets. In *AAAI'12 Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, 45–51.

[Hsueh, Melville, and Sindhwani 2009] Hsueh, P.-Y.; Melville, P.; and Sindhwani, V. 2009. Data quality from crowdsourcing: A study of annotation selection criteria. In *Proceedings of the NAACL HLT 2009 Workshop on Active Learning for Natural Language Processing*, 27–35.

[Karger, Oh, and Shah 2013] Karger, D. R.; Oh, S.; and Shah, D. 2013. Efficient crowdsourcing for multi-class labeling. In *Proceedings of the ACM SIGMETRICS/International Conference on Measurement and Modeling of Computer Systems, Pittsburgh, PA, USA*, 81–92.

[Kolobov, Mausam, and Weld 2013] Kolobov, A.; Mausam; and Weld, D. S. 2013. Joint crowdsourcing of multiple tasks. In *1st AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2013)*, 36–37.

[Lin, Mausam, and Weld 2015] Lin, C. H.; Mausam, M.; and Weld, D. S. 2015. Reactive learning: Actively trading off larger noisier training sets against smaller cleaner ones. In *Proceedings of the 32nd International Conference on Machine Learning, Lille, France (ICML)*. 5 pages.

[Rajpal, Goel, and Mausam 2015] Rajpal, S.; Goel, K.; and Mausam, M. 2015. POMDP-basedworker pool selection for crowdsourcing. In *Proceedings of the 32nd International Conference on Machine Learning, Lille, France (ICML)*. 9 pages.

[Riffe, Lacy, and Fico 2014] Riffe, D.; Lacy, S.; and Fico, F. 2014. *Analyzing media messages: Using quantitative content analysis in research*. Routledge, New York, N.Y.

[Tran-Thanh et al. 2013] Tran-Thanh, L.; Venanzi, M.; Rogers, A.; and Jennings, N. R. 2013. Efficient budget allocation with accuracy guarantees for crowdsourcing classification tasks. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, 901–908.

[Tumasjan et al. 2010] Tumasjan, A.; Sprenger, T. O.; Sandner, P. G.; and Welpe, I. M. 2010. Predicting elections with Twitter: What 140 characters reveal about political sentiment. In *Fourth International AAAI Conference on Weblogs and Social Media*, ICWSM 2010, 178–185.

[Wang et al. 2012] Wang, H.; Can, D.; Kazemzadeh, A.; Bar, F.; and Narayanan, S. 2012. A system for real-time twitter sentiment analysis of 2012 U.S. presidential election cycle. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics, Jeju, Republic of Korea*, 115–120.

[Yan et al. 2011] Yan, Y.; Rosales, R.; Fung, G.; and Dy, J. G. 2011. Active learning from crowds. In *Proceedings of the 28th International Conference on Machine Learning, Bellvue, WA*. 8 pages.

[Young and Soroka 2012] Young, L., and Soroka, S. 2012. Affective news: The automated coding of sentiment in political texts. *Political Communication* 29(2):205–231.